

AMENDMENTS TO THE CLAIMS

This listing of claims will replace all prior versions, and listings, of claims in the application:

- 1-22. (Canceled)
23. (Currently Amended) A machine-implemented method comprising the steps of:
 - dynamically selecting which occurrence counting technique to use from a plurality of available occurrence counting techniques by performing the steps of:
 - generating cost estimates for each of the plurality of available occurrence counting techniques based on an estimated I/O cost of using the available occurrence counting technique,
 - wherein generating cost estimates comprises performing:
 - determining a size of a candidate prefix tree;
 - determining an amount of memory that can be used for the candidate prefix tree;
 - comparing the size of the candidate prefix tree to the amount of memory that can be used to store the candidate prefix tree; and
 - generating an I/O cost estimate for the prefix tree technique based, at least in part, on the size of the candidate prefix tree and the amount of memory that can be used to store the candidate prefix tree; and
 - selecting the occurrence counting technique that has the lowest cost estimate; and
 - during a frequent itemset operation, using said selected occurrence counting technique to count occurrences of at least one combination to determine whether said at least one combination satisfies frequency criteria associated with said frequent itemset operation;

wherein the steps of dynamically selecting and using said selected occurrence counting technique are performed by one or more computing devices.

24. (Previously Presented) The machine-implemented method of Claim 23, wherein the selected occurrence counting technique is a prefix tree technique.

25. (Canceled)

26. (Previously Presented) The machine-implemented method of Claim 23, wherein the selected occurrence counting technique is a bitmap intersection technique.

27. (Previously Presented) The machine-implemented method of Claim 23, wherein generating cost estimates for each of the plurality of available occurrence counting techniques based on an estimated I/O cost comprises: generating an I/O cost estimate for a bitmap intersection technique based, at least in part, on a cost of reading bitmaps for each frequent item.

28. (Previously Presented) The machine-implemented method of Claim 23, wherein the plurality of available occurrence counting techniques include a bitmap intersection technique and a prefix tree technique.

29. (Previously Presented) The machine-implemented method of Claim 23, further comprising:
determining that a particular occurrence counting technique will not be considered during any phase of the frequent itemset operation; and
performing the frequent itemset operation without performing startup operations for said particular occurrence counting technique.

30-37. (Canceled)

38. (Currently Amended) A machine-implemented method comprising the steps of: dynamically selecting which occurrence counting technique to use from a plurality of available occurrence counting techniques based on conditions existing before a frequent itemset operation is performed in a computing environment in which the frequent itemset operation is to be performed,

wherein the conditions include how busy a computer system in which the frequent itemset operation is to be performed currently is; and

during said frequent itemset operation, using said selected occurrence counting technique to count occurrences of at least one combination to determine whether said at least one combination satisfies frequency criteria associated with said frequent itemset operation;

wherein the steps of dynamically selecting and using said selected occurrence counting technique are performed by one or more computing devices; wherein:

the frequent itemset operation is performed in a plurality of phases, wherein each phase is associated with combinations that have a particular number of items;

the step of dynamically selecting includes dynamically selecting which occurrence counting technique to use for at least one phase of said plurality of phases; and

the step of using includes using said selected occurrence counting technique to determine whether candidate combinations for said at least one phase satisfy said frequency criteria;

said at least one phase is a phase during which combinations having N items are processed;

a first occurrence counting technique is selected for said phase of said frequent itemset operation;

dynamically selecting a second occurrence counting technique in the phase of a subsequent frequent itemset operation during which combinations having N items are processed; and

wherein the first occurrence counting technique is different from said second occurrence counting technique.

39. (Previously Presented) The machine-implemented method of Claim 38, further comprising:
determining that a particular occurrence counting technique will not be considered during any phase of the frequent itemset operation; and

performing the frequent itemset operation without performing startup operations for said particular occurrence counting technique.

40. (Previously Presented) The machine-implemented method of Claim 38, wherein:
the plurality of phases comprises at least a first phase and a second phase;
the first phase is associated with combinations that have a first number of items;
the second phase is associated with combinations that have a second number of items; and
the occurrence counting technique selected for the first phase and the occurrence counting technique selected for the second phase are different.

41. (Previously Presented) A volatile or non-volatile computer-readable storage medium storing one or more sequences of instruction, wherein execution of the one or more sequences of instruction by one or more processors causes the one or more processors to perform:
dynamically selecting which occurrence counting technique to use from a plurality of available occurrence counting techniques by performing the steps of:
generating cost estimates for each of the plurality of available occurrence counting techniques based on an estimated I/O cost of using the available occurrence counting technique,
wherein generating cost estimates comprises performing:
determining a size of a candidate prefix tree;
determining an amount of memory that can be used for the candidate prefix tree;
comparing the size of the candidate prefix tree to the amount of memory that can be used to store the candidate prefix tree; and
generating an I/O cost estimate for a prefix tree technique based, at least in part, on the size of the candidate prefix tree and the amount of memory that can be used to store the candidate prefix tree;

selecting the occurrence counting technique that has the lowest cost estimate; and

during a frequent itemset operation, using said selected occurrence counting technique to count occurrences of at least one combination to determine whether said at least one combination satisfies frequency criteria associated with said frequent itemset operation.

42. (Previously Presented) A volatile or non-volatile computer-readable storage medium as recited in Claim 41, wherein the selected occurrence counting technique is a prefix tree technique.

43. (Previously Presented) A volatile or non-volatile computer-readable storage medium as recited in Claim 41, wherein the selected occurrence counting technique is a bitmap intersection technique.

44. (Previously Presented) A volatile or non-volatile computer-readable storage medium as recited in Claim 41, wherein generating cost estimates for each of the plurality of available occurrence counting techniques based on an estimated I/O cost comprises:

generating an I/O cost estimate for a bitmap intersection technique based, at least in part, on a cost of reading bitmaps for each frequent item.

45. (Previously Presented) A volatile or non-volatile computer-readable storage medium as recited in Claim 41, wherein the plurality of available occurrence counting techniques include a bitmap intersection technique and a prefix tree technique.

46. (Previously Presented) A volatile or non-volatile computer-readable storage medium as recited in Claim 41, wherein execution of the one or more sequences of instruction further causes the one or more processors to perform:

determining that a particular occurrence counting technique will not be considered during any phase of the frequent itemset operation; and

performing the frequent itemset operation without performing startup operations for said particular occurrence counting technique.

47. (Canceled)

48. (Previously Presented) A volatile or non-volatile computer-readable storage medium storing one or more sequences of instruction, wherein execution of the one or more sequences of instruction by one or more processors causes the one or more processors to perform:

dynamically selecting which occurrence counting technique to use from a plurality of available occurrence counting techniques based on conditions existing before a frequent itemset operation is performed in a computing environment in which the frequent itemset operation is to be performed,

wherein the conditions include how busy a computer system in which the frequent itemset operation is to be performed currently is, and an amount of volatile memory available to store a candidate prefix tree; and

during said frequent itemset operation, using said selected occurrence counting technique to count occurrences of at least one combination to determine whether said at least one combination satisfies frequency criteria associated with said frequent itemset operation;

wherein:

the frequent itemset operation is performed in a plurality of phases, wherein each phase is associated with combinations that have a particular number of items;

the dynamically selecting includes dynamically selecting which occurrence counting technique to use for at least one phase of said plurality of phases; and

the using includes using said selected occurrence counting technique to determine whether candidate combinations for said at least one phase satisfy said frequency criteria;

said at least one phase is a phase during which combinations having N items are processed;

a first occurrence counting technique is selected for said phase of said frequent itemset operation;

dynamically selecting a second occurrence counting technique in the phase of a subsequent frequent itemset operation during which combinations having N items are processed; and

wherein the first occurrence counting technique is different from said second occurrence counting technique.

49. (Previously Presented) A volatile or non-volatile computer-readable storage medium as recited in Claim 48, wherein execution of the one or more sequences of instruction further causes the one or more processors to perform:
determining that a particular occurrence counting technique will not be considered during any phase of the frequent itemset operation; and
performing the frequent itemset operation without performing startup operations for said particular occurrence counting technique.
50. (Previously Presented) A volatile or non-volatile computer-readable storage medium as recited in Claim 48, wherein:
the plurality of phases comprises at least a first phase and a second phase;
the first phase is associated with combinations that have a first number of items;

the second phase is associated with combinations that have a second number of items; and

the occurrence counting technique selected for the first phase and the occurrence counting technique selected for the second phase are different.